Robotics Research Technical Report

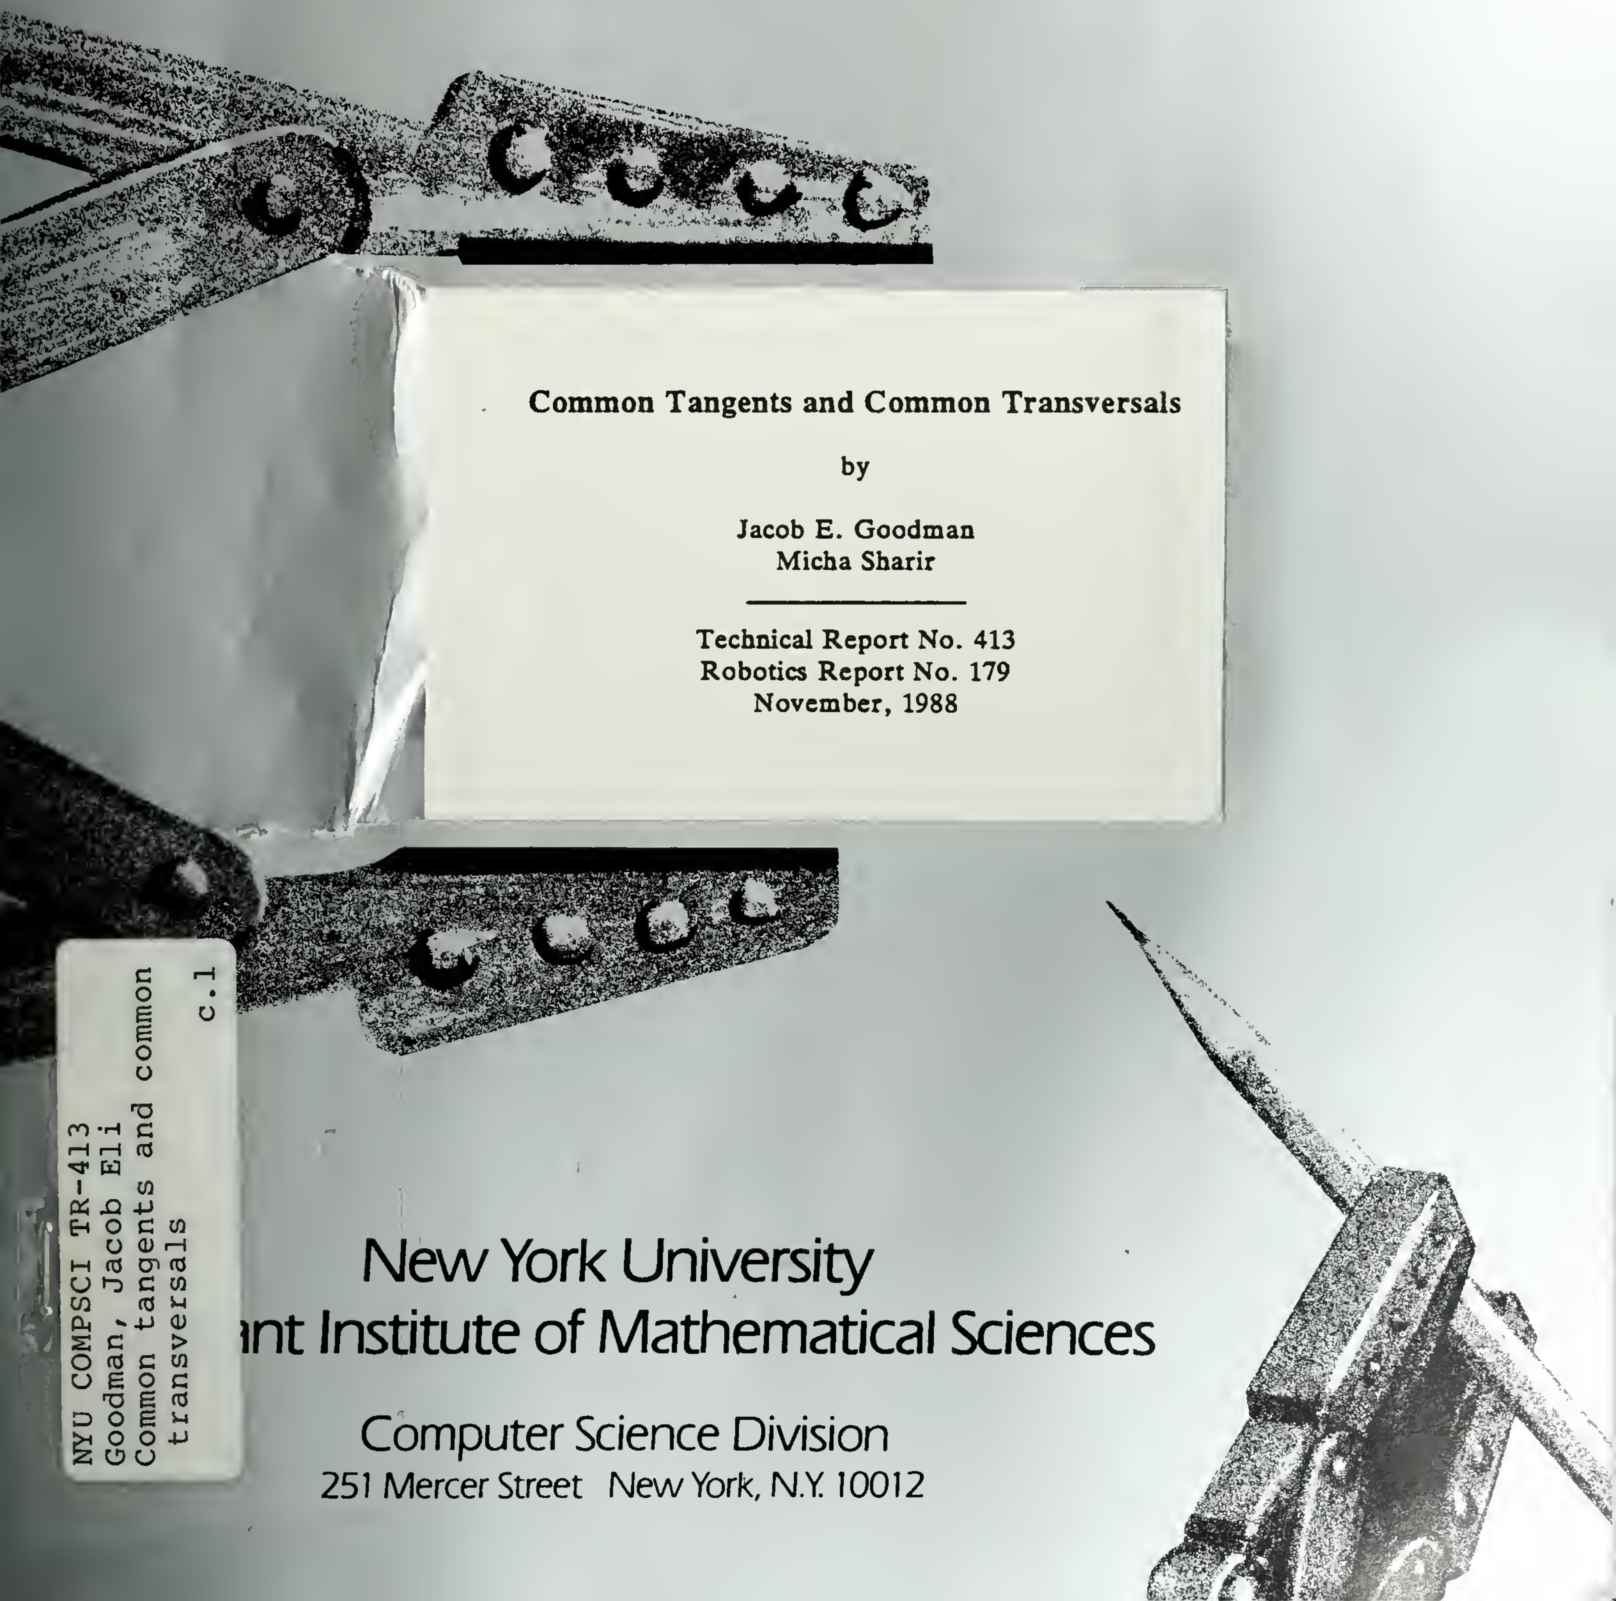

Common Tangents and Common Transversals

by

Jacob E. Goodman
Micha Sharir

Technical Report No. 413
Robotics Report No. 179
November, 1988

NYU COMPSCI TR-413
Goodman, Jacob Eli
Common tangents and common
transversals
c.1

New York University
nt Institute of Mathematical Sciences

Computer Science Division
251 Mercer Street New York, N.Y. 10012

Common Tangents and Common Transversals

by

Jacob E. Goodman
Micha Sharir

Technical Report No. 413
Robotics Report No. 179
November, 1988

New York University
Dept. of Computer Science
Courant Institute of Mathematical Sciences
251 Mercer Street
New York, New York 10012

Work by the first author has been supported in part by NSF grant DMS-85-01492 and PSC-CUNY grant 666426. Work by the second author has been supported in part by ONR grant N00014-87-K-0129, NSF grant DCR-83-20085, and grants from the Digital Equipment Corporation, the IBM Corporation, and the U.S.-Israeli Binational Science Foundation

Abstract

The paper consists of two parts. The first part shows that, given a family of $k \leq d$ compact convex and separated sets in E^d (i.e. no $(k-2)$-flat intersects all these sets), the space of all oriented common tangent hyperplanes to these sets, which separate them in any fixed manner (i.e. contain some specified subset of them on their positive side and the rest on their negative side) is homeomorphic to the sphere S^{d-k}. The second part exploits this result to obtain sharp bounds on the combinatorial complexity of the space of all common transversal hyperplanes to a given d-separated family of n compact convex sets in E^d, showing this complexity to be $O(n^{d-1})$.

Common Tangents and Common Transversals

Jacob E. Goodman*
Micha Sharir**

1. Introduction.

The purpose of this paper is twofold: to prove that under certain circumstances the space of common tangent hyperplanes to $k \leq d$ compact convex sets in $\mathbb{R}^d$ has a simple topological structure, and to show that this fact can be used to bound the combinatorial complexity of the set of hyperplanes which can be common transversals of a large family of separated convex sets. This then allows us, in turn, to derive an efficient algorithm for determining the space of common transversals if the convex sets are polytopal or, more generally, have sufficiently "simple" boundaries.

If K_1, K_2, and K_3 are three disjoint spheres in "general position" in $\mathbb{R}^3$ (i.e., not all met by any line), and $\mathcal{T}$ is the family of common tangent planes separating them in a given way, say the first two from the third, or all three on one side, then it is not difficult to see that $\mathcal{T}$ consists of precisely two members. Similarly, if K_1 and K_2 are two disjoint spheres in $\mathbb{R}^3$ and $\mathcal{T}$ is the family of planes tangent to K_1 and K_2, say "internally", it is clear that $\mathcal{T}$ consists of a circle of tangent planes. If we clarify the situation by considering instead *oriented* planes H and assigning to each of K_1, K_2 a sign indicating on which side of H it is to lie, this circle, or the pair of points in the first case, has a natural embedding in the sphere S^2 of all unit normal vectors to oriented planes in $\mathbb{R}^3$.

In this paper we consider to begin with the more general problem in which the dimension of the ambient space is any number $d \geq 2$, and the spheres are replaced by $k \leq d$ suitably separated compact convex sets. Our result (Theorems 1 and 2 below) is that if $K_1, \ldots, K_k$

---

*Supported in part by NSF grant DMS-85-01492 and PSC-CUNY grant 666426.
**Supported in part by ONR grant N00014-87-K-0129, NSF grant DCR-83-20085, and grants from the Digital Equipment Corporation, the IBM Corporation, and the U.S.-Israeli Binational Science Foundation.

are separated, then the space of common supporting oriented hyperplanes dividing them in a given way is homeomorphic to a sphere of dimension $d - k$.

Recall [**5**] that a family $\mathbf{K} = \{K_1, \dots, K_d\}$ of convex sets is said to be *separated* in R^d if it has no $(d-2)$-transversal (i.e., if there is no $(d-2)$-flat meeting all the members of $\mathbf{K}$), or equivalently if any r of the sets can be separated from the remaining $d - r$ by a hyperplane; here we shall extend this to a family $\mathbf{K}$ of $k \leq d+1$ sets by requiring that the members of $\mathbf{K}$ have no $(k-2)$-transversal. (Thus two sets are separated if they have no point in common, three if they are met by no line, and so on.) It turns out that this is precisely the notion of "general position" needed to generalize disjointness (in the case of two convex sets) which permits us to extend the results mentioned above for spheres to convex sets in higher dimensions.

Surprisingly, the only work relating to this result which appears in the literature seems to be a paper by H. Brunn [**2**], which is restricted to the case of two convex bodies in R^3, and from which our results can easily be seen to follow in that special low-dimensional case.

In Section 3 we apply the results of Section 2 on the topology of the space of common tangent hyperplanes to derive an asymptotic bound on the complexity of the space $\tau(\mathbf{K})$ of common *transversal* hyperplanes to a collection $\mathbf{K}$ of n d-wise separated compact convex sets in R^d, for n large. This problem has been considered previously in the computer science literature, but only in the special cases of dimension 2 (see [**4**] for example, where a tight bound of $2n - 2$ was obtained on the complexity of $\tau(\mathbf{K})$) and dimension 3 (see [**3**] and Remark (3) in Section 4 below).

Here we are able to show (Theorem 3) that for $d \geq 2$ the complexity of $\tau(\mathbf{K})$ is $O(n^{d-1})$, which extends the results previously known in the lower-dimensional cases. In particular, we show that this bound applies as well to the number of distinct order types in which a transversal may meet the members of a given family of separated sets, a problem which (in dimension 2 at least) has drawn some recent attention ([**4**], [**7**], [**11**]).

Finally, in Section 4, we present an algorithm which uses the analysis of the preceding sections to calculate efficiently the complete space of hyperplanes transversal to a given family of separated compact convex sets in $\mathbb{R}^d$, provided the sets themselves are described in a way which permits "constant-time" manipulation of any fixed number of them, e.g., if each set is a convex polytope with a bounded number of sides. If there are n sets, the space of common transversals can be calculated in time $O(n^{d-1}\log^2 n)$.

Section 4 concludes with a few remarks on open problems.

We would like to express our gratitude to János Pach for some helpful discussions, and to acknowledge that Tibor Bisztriczky has independently found a proof of Theorem 2 in the special case $k = d$ [**1**], along somewhat different lines from ours.

2. The space of common tangents.

LEMMA 1. *Let* $\mathbf{K} = \{K_0, K_1, \ldots, K_d\}$ *be a separated family of compact convex sets in* $\mathbb{R}^d$ *with* K_0 *a point, and let* $\mathbf{K} = \mathbf{P} \cup \mathbf{N}$ *be a partition of* $\mathbf{K}$ *into two disjoint subsets. Suppose there is an oriented hyperplane* H *which has the members of* $\mathbf{P}$ *on its positive side and those of* $\mathbf{N}$ *on its negative side, and which supports* $K_1, \ldots, K_d$*. Then* H *is unique.*

PROOF: Suppose H and H' are two such oriented hyperplanes. Let F be the $(d-2)$-flat in which they meet, and assume to begin with that F is at finite distance. Together, H and H' divide $\mathbb{R}^d$ into four quadrants, of which one, call it Q_1, lies on the positive side of both H and H', and the opposite one, say Q_2, on the negative side of both. It is clear that the K_i in $\mathbf{P}$ lie in Q_1, while those in $\mathbf{N}$ lie in Q_2. Rotating H about F through Q_1 and Q_2 simultaneously until it coincides with H', H passes through all of $K_1, \ldots, K_d$ at every stage, hence never through K_0 (by the separation condition). But at the end, all of $K_1, \ldots, K_d$ lie on the *wrong* side of the rotated H, hence—reversing the orientation to get H'—we see that H' cannot have K_0 on the correct side.

If H and H' are parallel, we need only move H parallel to itself until it coincides with H', and then the rest of the proof goes through. □

LEMMA 2 (PROJECTION LEMMA). *If $K_1, \ldots, K_k$ form a separated family of $k \leq d$ compact convex sets in R^d, then so do their orthogonal projections into some $(k-1)$-dimensional subspace of R^d.*

PROOF: Enlarging each K_i slightly if necessary, we may assume that the sets K_i are strictly convex and have non-empty interior.

We begin by proving the result first in the case where $k = d$, and then indicate the modifications necessary for $k < d$. For this purpose, we use some of the machinery developed in [**5**], which, in turn, was inspired by [**6**]. For each point x at infinity, i.e., each direction of parallel projection, let π_x be the projection mapping of R^d onto the hyperplane H_x normal to the direction represented by x and passing through the origin. We seek a point x at infinity such that $\pi_x(K_1), \ldots, \pi_x(K_d)$ are separated, i.e., such that x is not co-hyperplanar with any choice of points $y_1, \ldots, y_d$, $y_i \in K_i$.

Suppose there is no such point x. Fix an interior point $z_i \in K_i$ for each i, and—for every $t \in [0,1]$—let $K_i(t)$ be the contraction of K_i by the factor t about z_i. Since $K_1, \ldots, K_d$ are separated, the points $z_1, \ldots, z_d$ are in general position, hence we can certainly find a point $x(0)$ at infinity not co-hyperplanar with $K_1(0), \ldots, K_d(0)$. As t increases from 0 to 1, it gets progressively harder to find a point $x(t)$ not co-hyperplanar with $K_1(t), \ldots, K_d(t)$, i.e., the values of t for which this is possible form an interval of the form $[0, \tau)$, $\tau < 1$, which in fact must be open at the right since any projection which keeps the images of the sets $K_i(t)$ separated does the same for the sets $K_i(t+\epsilon)$. Since the hyperplane at infinity is compact, we can find an increasing sequence $t_1, t_2, \ldots$ converging to τ and corresponding points $x_1, x_2, \ldots$ converging to a point x such that x_n is not co-hyperplanar with the sets $K_i(t_n)$ for each $n = 1, 2, \ldots$. Further, by rotating R^d appropriately for each n, we may assume that $x_n = x$ for each n, i.e., that we have a sequence $K_1^n, \ldots, K_d^n$ of families of sets converging (in the Hausdorff metric) to $K_1(\tau), \ldots, K_d(\tau)$, such that for each n the corresponding sets are not co-hyperplanar with x, while in the limit they are.

Since π_x does not leave the sets $K_i(\tau)$ separated, there is a hyperplane H through x which meets all the sets $K_i(\tau)$. It is clear that H must support them—otherwise we could rotate H slightly to a hyperplane H' still passing through x and cutting through the interiors of *all* of the sets $K_i(\tau)$, and then we could find a family $K_1^n, \ldots, K_d^n$ sufficiently close which H' would still meet, which is impossible. Let $y_i = H \cap K_i(\tau)$ for each $i = 1, \ldots, d$. We claim that their projections $\pi_x(y_i)$ on H_x form a minimal Radon partition on $H \cap H_x$ [5], i.e., that those corresponding to K_i's lying on one side of H cannot be separated from those corresponding to K_i's lying on the other side. This follows as in [5, p. 306]; in brief, the reason is that otherwise we could rotate H slightly to cut through the interior of all of the sets $K_i(\tau)$, which we have seen above is impossible. Since there are fewer than $d+1$ points of contact y_i, we can rotate H slightly to move it *away* simultaneously from all of the sets $K_i(\tau)$, so that some appropriate projection $\pi_{x'}$ (with x' close to x) moves the sets $K_i(\tau)$ away from $H \cap H_x$ on both sides. But then it follows from Lemma 3 of [5] that *no other* $(d-2)$-flat can meet the sets $\pi_{x'}(K_1(\tau)), \ldots, \pi_{x'}(K_d(\tau))$, which gives a contradiction.

If $k < d$, we make the following modifications: x should now be thought of as a $(d-k)$-flat at infinity, and π_x as the corresponding projection onto the $(k-1)$-space orthogonal to the $(d-k+1)$-space spanned by x. It is now the compactness of the Grassmannian $G_{d-k+1,d}$ of subspaces of $\mathbf{R}^d$ of dimension $d-k+1$ which we invoke to find our convergent sequences, and the rest of the details remain as above. ☐

THEOREM 1. *Let* $\mathbf{K} = \{K_0, K_1, \ldots, K_d\}$ *be a separated family of compact convex sets in* $\mathbf{R}^d$ *with* K_0 *a point, and let* $\mathbf{K} = \mathbf{P} \cup \mathbf{N}$ *be a partition of* $\mathbf{K}$ *into two disjoint subsets. Then there exists one and only one oriented hyperplane* H *which has the members of* $\mathbf{P}$ *on its positive side and those of* $\mathbf{N}$ *on its negative side, and which supports* $K_1, \ldots, K_d$*. Moreover,* H *varies continuously with* $K_0, \ldots, K_d$ *in the Hausdorff topology, as long as these sets remain separated.*

PROOF: The uniqueness has already been proved in Lemma 1. For the existence we proceed by induction on d, the assertion being clear for $d = 1$. Assuming $d \geq 2$, let

$\pi_x : \mathbb{R}^d \to \mathbb{R}^{d-1}$ be a separating projection for $K_1, \ldots, K_d$, as in Lemma 2; i.e., the images $\pi_x(K_1), \ldots, \pi_x(K_d)$ in $\mathbb{R}^{d-1}$ have no transversal $(d-2)$-flat. By induction hypothesis, there is a unique oriented $(d-2)$-flat F in $\mathbb{R}^{d-1}$ separating those sets (with their inherited signs) correctly, and supporting $\pi_x(K_1), \ldots, \pi_x(K_{d-1})$. (Notice that since $\pi_x(K_1), \ldots, \pi_x(K_d)$ are separated, it makes no difference which point of $\pi_x(K_d)$ we choose to play the role of the "K_0" in the induction hypothesis.) Also, again by induction hypothesis, there is a second oriented $(d-2)$-flat F' with the same properties as F, except that it has $\pi_x(K_d)$ on its "wrong" side. Lifting these up to $\mathbb{R}^d$ by π_x^{-1}, we get oriented hyperplanes H and H' each supporting $K_1, \ldots, K_{d-1}$, with H separating $K_1, \ldots, K_d$ correctly, and H' separating $K_1, \ldots, K_{d-1}$ correctly but having K_d on its wrong side.

Again by Lemma 2, there is a separating projection $\pi_y : \mathbb{R}^{d-1} \to \mathbb{R}^{d-2}$ for the sets $\pi_x(K_1), \ldots, \pi_x(K_{d-1})$. This implies that the map

$$\pi_L = \pi_y \circ \pi_x : \mathbb{R}^d \to \mathbb{R}^{d-2}$$

is a separating projection for the family $\{K_1, \ldots, K_{d-1}\}$, where $L = \pi_x^{-1}(y)$, so that as z varies continuously in L we get a continuously varying circle of projections π_z onto an $\mathbb{R}^{d-1}$ (which we may regard as remaining fixed), such that each family $\{\pi_z(K_1), \ldots, \pi_z(K_{d-1})\}$ is separated. These projected sets $\{\pi_z(K_i)\}$ vary continuously with z, so by induction hypothesis the corresponding supporting hyperplanes F_z in $\mathbb{R}^{d-1}$ also vary continuously; hence as their pullback $H_z = \pi_z^{-1}(F_z)$ to $\mathbb{R}^d$ moves from H to H' (which corresponds to z moving along L from x to "$-x$") there must be at least one position H_1 in which H_z separates all of $K_1, \ldots, K_d$ correctly, and supports K_d as well, and a second as z continues along L back to x again, say H_2, in which H_z does the same. By Lemma 1, however, at most one of these has K_0 on its positive side, and at most one on its negative side. Hence *at least* one must have K_0 on its correct side.

Finally, we can see that the common tangent hyperplane H varies continuously with $K_0, \ldots, K_d$ as follows: For each oriented hyperplane H with positive normal vector $\mathbf{u}$

having K_0 on its correct side, let $d_{H,i}$ be the distance H must be moved in a normal direction to support K_i on the correct side, and let

$$D_H = \sum_{i=1}^{d} d_{H,i}^2.$$

Then we have been seeking, and have found, a unique hyperplane H which minimizes the continuous function D_H over an appropriate compact subset of the space of all hyperplanes in R^d. (In fact, this minimum has turned out to be 0!) But since the function D_H itself depends continuously on $K_0, \ldots, K_d$ as well as on H, so does the hyperplane achieving the minimum, by the uniqueness and the compactness. This concludes the proof. $\square$

REMARK: It is also possible to prove Theorem 1 along the lines suggested in the last paragraph: find a hyperplane H which minimizes the function D_H, and prove it must satisfy the requirements of the theorem by rotating it slightly to reduce D_H if it does not; the details are somewhat more intricate than those above.

Notice also that if several of the K_i happen to be points, even points at infinity, the conclusion of Theorem 1 remains valid, since we can always move them to finite distance by a permissible projective transformation, i.e., we can move the hyperplane at infinity away from them without affecting the remaining convex bodies in any material way. (Of course if some K_i, $1 \leq i \leq d$, is a point, the "positive" and "negative" side condition for it is vacuous.)

THEOREM 2. *Let* $\mathbf{K} = \{K_1, \ldots, K_k\}$, $k \leq d$, *be a separated family of compact convex sets in* R^d, *and let* $\mathbf{K} = \mathbf{P} \cup \mathbf{N}$ *be a partition of* $\mathbf{K}$ *into two disjoint subsets. Let* T *be the set of oriented hyperplanes* H *which have the members of* $\mathbf{P}$ *on their positive side and those of* $\mathbf{N}$ *on their negative side, and which support* $K_1, \ldots, K_k$. *Then* T *(with its members identified with their positive normals in* S^{d-1}*) is a topological sphere of dimension* $d - k$.

PROOF: Theorem 1 disposes of the case $k = d$; let us therefore suppose that $k < d$. By Lemma 2 we can find a $(d-k)$-flat F at infinity such that for any choice of $x_1 \in K_1, \ldots, x_k \in K_k$, the x_i and F are in general position.

Consider a $(d-k-1)$-flat G lying in F, choose an orientation for G, i.e. a "positive" side of G in F, and choose $d-k$ points $P_{k+1}, \ldots, P_d$ in general position in G. By Theorem 1 and the Remark following, there is one and only one oriented hyperplane H supporting $K_1, \ldots, K_k$ correctly, passing through $P_{k+1}, \ldots, P_d$ (hence cutting F in G), and determining the same orientation in F as does G. (Intuitively, H is the unique member of T whose normal is orthogonal to G and points to the positive side of G.) If we reverse the orientation of G then the oriented hyperplane H changes, while if we move G to a new position H also moves, since otherwise H would meet F in two distinct $(d-k-1)$-flats and would therefore contain F entirely, contrary to the general position assumption.

On the other hand, any hyperplane H which supports $K_1, \ldots, K_k$ and separates them correctly must cut F in *some* $(d-k-1)$-flat G, so there is a bijection f from the set of oriented $(d-k-1)$-flats contained in F into the set T of hyperplanes supporting $K_1, \ldots, K_k$ and separating them correctly. Moreover, f is continuous by Theorem 1. Since the $(d-k-1)$-flats contained in F form a projective space P^{d-k} of dimension $d-k$, and the specification of an orientation for each flat is just the lift of P^{d-k} to its covering space S^{d-k} (recall that P^{d-k} is obtained by identifying antipodal points of S^{d-k}), the conclusion follows. ☐

REMARK: If $\mathbf{K}$ consists simply of k points in general position, the sphere T given by Theorem 2 reduces to the intersection of S^{d-1} with the $(d-k+1)$-space normal to the affine span of $\mathbf{K}$. By deforming the sets K_i in Theorem 2 continuously to points, and applying the theorem at each stage, one can then show that the pair (S^{d-1}, T) is isotopic to the standard pair (S^{d-1}, S^{d-k}), so that—in particular—T is always unknotted in the ambient sphere S^{d-1}.

3. The space of common transversals.

Let $\mathbf{K} = \{K_1, \ldots, K_n\}$ be a collection of $n >> d$ compact convex sets in $\mathbf{R}^d$, any d of which are separated. A hyperplane H is a *common transversal* (or a *stabbing hyperplane*) of

$\mathbf{K}$ if $H \cap K_i \neq \emptyset$ for all $i = 1, \ldots, n$. In this section we apply the previous results to analyze the combinatorial structure of the space $\tau(\mathbf{K})$ of all common transversal hyperplanes of $\mathbf{K}$, and in particular to bound its complexity.

To define $\tau(\mathbf{K})$ more precisely, we represent each oriented hyperplane H by a pair $(\mathbf{u}, d)$, where $\mathbf{u} \in S^{d-1}$ is a unit vector in the direction of the normal to H, and d is the *signed* distance from the origin to H, i.e., $d = \langle a, \mathbf{u} \rangle$ for any point a on H; notice that $(\mathbf{u}, d)$ and $(-\mathbf{u}, -d)$ are the two oriented representations of the same unoriented hyperplane. The space $\tau(\mathbf{K})$ is thus the subset of $S^{d-1} \times \mathsf{R}$ consisting of all pairs $(\mathbf{u}, d)$ representing the common transversal hyperplanes to $\mathbf{K}$.

To get a better handle on the structure of $\tau(\mathbf{K})$, we define, for each $i = 1, \ldots, n$, a function f_i on the unit sphere S^{d-1} as follows. For each vector $\mathbf{u} \in S^{d-1}$, let $f_i(\mathbf{u})$ be the signed distance from the origin to the hyperplane supporting K_i and having positive normal $\mathbf{u}$. Let

$$f(\mathbf{u}) = \min_i f_i(\mathbf{u}), \quad \mathbf{u} \in S^{d-1}.$$

Let H be a common transversal hyperplane of $\mathbf{K}$ and choose an orientation so that H has some normal direction $\mathbf{u}$. Then the signed distance d of H from the origin (measured, as above, by $\langle a, \mathbf{u} \rangle$, where a is any point on H) must satisfy $d \leq f_i(\mathbf{u})$ for all i, or $d \leq f(\mathbf{u})$. Similarly, if we take $-\mathbf{u}$ as the normal direction to H, we obtain $-d \leq f(-\mathbf{u})$. Thus we must have

$$f(\mathbf{u}) + f(-\mathbf{u}) \geq 0.$$

It is also easily checked that any $\mathbf{u} \in S^{d-1}$ that satisfies this inequality admits a common transversal hyperplane to $\mathbf{K}$ orthogonal to $\mathbf{u}$. Furthermore, the desired space $\tau(\mathbf{K})$ can be represented as

$$\tau(\mathbf{K}) = \{(\mathbf{u}, d) : \mathbf{u} \in S^{d-1},\ d \in \mathsf{R},\ f(\mathbf{u}) + f(-\mathbf{u}) \geq 0,\ -f(-\mathbf{u}) \leq d \leq f(\mathbf{u})\}.$$

In particular, complete information about $\tau(\mathbf{K})$ is provided by its projection $\sigma(\mathbf{K})$ on

S^{d-1}, i.e.,

$$\sigma(\mathbf{K}) = \{\mathbf{u} \in S^{d-1} : f(\mathbf{u}) + f(-\mathbf{u}) \geq 0\},$$

together with a pair of partitions of this space, denoted by $\sigma^+(\mathbf{K})$ and $\sigma^-(\mathbf{K})$, into maximal connected cells of dimensions between 0 and $d-1$, so that to each cell in $\sigma^+(\mathbf{K})$ (resp. $\sigma^-(\mathbf{K})$) there corresponds a fixed known function $f_i(\mathbf{u})$ (resp. $f_j(-\mathbf{u})$) (or a fixed set of such functions) which attains the minimum $f(\mathbf{u})$ (resp. $f(-\mathbf{u})$). Our goal is to analyze the *combinatorial complexity* of $\sigma^+(\mathbf{K})$ and of $\sigma^-(\mathbf{K})$, namely the number of cells in these two partitions.

To analyze the combinatorial structure of $\sigma^+(\mathbf{K})$, say, we proceed as follows. Fix an index $i = 1, \ldots, n$. For each $j \neq i$ consider the following two cells on S^{d-1}:

$$R_j^+ = \{\mathbf{u} \in S^{d-1} : f_j(\mathbf{u}) < f_i(\mathbf{u})\}$$

$$R_j^- = \{\mathbf{u} \in S^{d-1} : f_j(-\mathbf{u}) + f_i(\mathbf{u}) < 0\}.$$

(Intuitively, R_j^+ consists of those directions in which a sweep-hyperplane leaves K_j before leaving K_i, while R_j^- consists of those directions in which it leaves K_i before entering K_j.) Let

$$R = R^{(i)} = \bigcup_{j \neq i} (R_j^+ \cup R_j^-).$$

Then clearly $\mathbf{u} \in S^{d-1} \setminus R$ if and only if $f(\mathbf{u}) + f(-\mathbf{u}) \geq 0$ and $f(\mathbf{u}) = f_i(\mathbf{u})$. In other words, the complement of R is exactly that portion of $\sigma^+(\mathbf{K})$ where the minimum, f, is attained by f_i. It thus suffices to determine an upper bound on the combinatorial complexity of the boundary of R; multiplying this bound by n will then yield a bound on the complexity of $\sigma^+(\mathbf{K})$.

LEMMA 3. *The set $\{\partial R_j^+\}_{j \neq i} \cup \{\partial R_j^-\}_{j \neq i}$ is an arrangement of pseudo-spheres within S^{d-1} (i.e., the intersection of any subcollection of $1 \leq k \leq d$ of these objects is homeomorphic to the unit sphere S^{d-k-1}, just as would be the case for an arrangement of great spheres in general position).*

PROOF: Let

$$\{\partial R^+_{i_1}, \ldots, \partial R^+_{i_p}, \partial R^-_{j_1}, \ldots, \partial R^-_{j_q}\}$$

be a subcollection of $k = p + q$ of these boundaries. It follows from the definition that their intersection is the subset of S^{d-1} consisting of the normal directions $\mathbf{u}$ of all oriented hyperplanes H such that H is tangent to $K_i, K_{i_1}, \ldots, K_{i_p}, K_{j_1}, \ldots, K_{j_q}$, and such that the first $p+1$ sets lie on the negative side of H and the last q on the positive side. By Theorem 2, this space is homeomorphic to S^{d-k-1}, as asserted. □

The following proposition follows immediately from the reults of [8] (see also [9] and [10]), but we include a sketch of the proof for the sake of completeness.

PROPOSITION 1. *The combinatorial complexity of the boundary of the union of n pseudo-hemispheres (i.e., sets whose boundaries constitute an arrangement of pseudo-spheres in the sense defined above) in S^{d-1}, for $d \geq 3$, is $O(n^{d-2})$.*

PROOF: We use induction on d. The base case $d = 3$ is merely a special case of the result of [8] that the union of n Jordan regions in the plane, bounded by Jordan curves with the property that each pair of them intersects at most twice (exactly twice in our case), has only a linear number of such intersections on its boundary. For $d > 3$ we apply the same idea as above: Let $B_1, \ldots, B_n$ be the given pseudo-hemispheres. We fix one of them, say B_1, and consider the intersections $C_i = \partial B_1 \cap B_i$, for $i \neq 1$. It is easily checked that the C_i's form a collection of pseudo-hemispheres within ∂B_1, which is homeomorphic to S^{d-2}. By induction hypothesis, the boundary of the union of the C_i's has combinatorial complexity $O(n^{d-3})$. But this boundary is plainly equal to the boundary of that portion of ∂B_1 which appears along the boundary of the union of the B_i's. Thus the total combinatorial complexity of the boundary of this union is at most $n \cdot O(n^{d-3}) = O(n^{d-2})$, as claimed. □

As a corollary, we immediately obtain our main result:

THEOREM 3. *The combinatorial complexity of $\sigma^+(\mathbf{K})$ and of $\sigma^-(\mathbf{K})$, hence also of $\tau(\mathbf{K})$, is $O(n^{d-1})$.*

4. Algorithms and remarks.

Assuming that the sets K_i have relatively simple shape so that manipulating any fixed number of the functions $f_i(\mathbf{u})$ can be assumed to take constant time, we can calculate $\tau(\mathbf{K})$ efficiently. The method follows the line of proof of Proposition 1, tracing the induction down to dimension 2, at which point it applies the algorithm given in [8]. The entire procedure takes time $O(n^{d-1}\log^2 n)$.

To avoid the rather tedious details of calculating the topological structure of $\tau(\mathbf{K})$, we restrict the algorithm to produce only the vertices of that set, i.e., only extreme common transversal hyperplanes that are also common tangents to d of the given sets.

ALGORITHM: We first describe the inner routine of the algorithm, which recurses on the dimension, and then give the top-level part.

```
Procedure UNION:
Input: A dimension $r \geq 2$ and a set of pseudo-hemispheres $B_1, \ldots, B_q$ on the sphere $S^r$.
Output: The vertices of $\bigcup_j B_j$.
if $r = 2$ then
    $B_1, \ldots, B_q$ are "pseudo-disks" on $S^2$.
    Apply the divide-and-conquer technique of [8] to calculate the desired vertices
        This assumes that the $B_j$'s have some reasonable monotonicity properties, to allow
        us to apply line sweeping. (This is the case, for example, if each $B_i$ is actually a
        hemisphere, which happens if all the original sets $K_i$ are balls.) The algorithm takes
        $O(q\log^2 q)$ time.
else
    for each $i = 1, \ldots, q$
        for each $j \neq i$ let $B_j^* := \partial B_i \cap B_j$
        call UNION with $r - 1$ and $\{B_j^* : j \neq i\}$
            The $B_j^*$ are pseudo-hemispheres on the $(r-1)$-sphere $\partial B_i$.
```

```
  endfor
  Return the union of the outputs of all recursive calls, with duplications removed
    Each vertex is produced $r$ times.
endif
end UNION

MAIN ROUTINE:
Input: The dimension $d$ and a family $\mathbf{K}$ of $n$ $d$-wise separated convex sets in $\mathbb{R}^d$.
Output: The collection of all extreme common transversal hyperplanes of $\mathbf{K}$.
for each $i = 1, \ldots, n$
  Form the pseudo-hemispheres $R_j^+$, $R_j^-$ for all $j \neq i$
  call UNION with $d-1$ and the set of these pseudo-hemispheres
endfor
Return the union of the outputs of all calls to UNION, with duplications removed
  Each vertex is produced $d$ times.
end
```

Every vertex is an extreme common transversal, tangent to those d sets K_i for which the top-level call to UNION has produced that vertex.

The time complexity of the algorithm is an immediate consequence of Proposition 1, the discussion preceding it, and the time complexity of the innermost stage of the algorithm, as borrowed from [8]. Notice that the only place in the algorithm where the actual shape of the sets K_i really matters is at the bottom of the recursion, at $r = 2$. In other words, we need to assume only that the shape of each pseudo-circle of oriented hyperplanes that are tangent to any fixed $d-1$ sets in $\mathbf{K}$ and separate them in any specific manner is simple enough for the algorithm of [8] to be applicable.

REMARKS:

(1) An open problem is whether the bound provided by Theorem 3 is tight in the worst

case, for separated convex sets. It is easy to construct an example where the complexity of $\tau(\mathbf{K})$ is $\Omega(n^{\lfloor \frac{d}{2} \rfloor})$, as follows. Assume for simplicity that d is a power of 2. For the base case $d = 2$ use the construction in [4]. For $d > 2$, take two complementary subspaces V_1, V_2 of R^d of dimension $d/2$ each. Construct inductively in each of them a family $\mathbf{K}_1$, $\mathbf{K}_2$ of $n/2$ separated convex sets whose "stabbing region" has combinatorial complexity $\Omega(n^{\frac{d}{4}})$, and observe that any pair of hyperplanes (in V_1, V_2 respectively) which are common transversals in $\mathbf{K}_1$, $\mathbf{K}_2$, respectively, span a hyperplane in R^d which is a common transversal to the separated family $\mathbf{K} = \mathbf{K}_1 \cup \mathbf{K}_2$. This implies easily that the combinatorial complexity of $\tau(\mathbf{K})$ is at least $\Omega((n^{\frac{d}{4}})^2) = \Omega(n^{\frac{d}{2}})$, the constant depending on d. It is easy to extend this construction to the case of general d. Can this lower bound be improved? Alternatively, can the upper bound of Theorem 3 be improved?

(2) It is instructive to notice that the analysis above also provides a means to study the combinatorial complexity of the convex hull of $\mathbf{K}$. Indeed, for this purpose we want to consider the space of all hyperplanes that contain all the sets in $\mathbf{K}$ on the same side. But the boundary of this space is easily seen to coincide with the *upper envelope* of the graphs of the functions $f_i(\mathbf{u})$, as defined above, rather than with their lower envelope f. Thus our analysis allows us to conclude that the combinatorial complexity of the convex hull of $\mathbf{K}$ is $O(n^{d-1})$. Generally speaking, there does not exist complete symmetry between these two problems, because the shape of the functions f_i seems to render the structure of their upper envelope simpler than that of their lower envelope. Nonetheless, this connection between convex hulls and common transversals is intriguing, and raises several open problems, the most obvious being to obtain sharp bounds on the complexity of the convex hull of n (d-wise separated) convex sets in R^d. If these sets reduce to points, the complexity of the hull is at most $O(n^{\lfloor \frac{d}{2} \rfloor})$, as implied by the "upper bound theorem" for convex polytopes. Is this also an upper bound for the case where the sets in $\mathbf{K}$ are larger than points (but still separated)? See [10] for a discussion of some special cases of the convex hull problem for sets in three dimensions.

(3) If the sets in **K** are not separated then in general very little can be said about the complexity of $\tau(\mathbf{K})$. For example, if two convex sets are allowed to overlap, even the space of all common tangent hyperplanes to both sets can be of arbitrarily large combinatorial complexity. In certain non-separated cases, however, sharp bounds can still be obtained for the complexity of $\tau(\mathbf{K})$. For example, if the sets in **K** are all convex polytopes in three dimensions (not necessarily separated), then it is shown in [3] that the combinatorial complexity of $\tau(\mathbf{K})$ is $O(n^2\alpha(n))$, where n is the total number of vertices of these polytopes and where $\alpha(n)$ is the extremely slowly growing inverse Ackermann function. The bound reduces to only $O(n^2)$ for a collection of n line segments in 3-space.

(4) In addition to the combinatorial complexity of $\tau(\mathbf{K})$, there are two other parameters of interest. One is the number $\nu(\mathbf{K})$ of connected components of $\tau(\mathbf{K})$, and the other is the number $\mu(\mathbf{K})$ of different order types that can be induced by common transversal hyperplanes of **K**, where the order type induced by a transversal hyperplane H in $\mathbb{R}^d$ is determined by choosing a point in each intersection $H \cap K_i$, $i = 1, \ldots, n$, and reading off the orientations of all d-tuples of points selected. Since the sets in **K** are separated, it is easily seen that the order type induced by H is uniquely defined, and does not change as H varies in a single connected component of $\tau(\mathbf{K})$, so that $\mu(\mathbf{K}) \leq \nu(\mathbf{K})$; and it is equally clear that $\nu(\mathbf{K})$ does not exceed the combinatorial complexity of $\tau(\mathbf{K})$. Thus the upper bound established in Theorem 3 applies to $\mu(\mathbf{K})$ and to $\nu(\mathbf{K})$ as well. In $\mathbb{R}^2$, moreover, the results of [4] show that the maximum possible values of all three parameters are the same (namely $2n - 2$ for n sets). Does a similar equality hold in higher dimensions?

REFERENCES

1. T. Bisztriczky, *On separated families of convex bodies*, preprint.
2. H. Brunn, *Sätze über zwei getrennte Eikörper*, Math. Annalen **104** (1931), 300–324.
3. H. Edelsbrunner, L. Guibas, and M. Sharir, *The upper envelope of piecewise linear functions: algorithms and applications*, Discrete Comput. Geom. (to appear).

4. H. Edelsbrunner and M. Sharir, *The maximum number of ways to stab n convex non-intersecting objects in the plane is* $2n-2$, Discrete Comput. Geom. (to appear).
5. J. E. Goodman and R. Pollack, *Hadwiger's transversal theorem in higher dimensions*, J. Amer. Math. Soc. **1** (1988), 301–309.
6. H. Hadwiger, *Über Eibereiche mit gemeinsamer Treffgeraden*, Portugal. Math. **16** (1957), 23–29.
7. M. Katchalski, T. Lewis, and J. Zaks, *Geometric permutations for convex sets*, Discrete Math. **54** (1985), 271–284.
8. K. Kedem, R. Livne, J. Pach, and M. Sharir, *On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles*, Discrete Comput. Geom. **1** (1986), 59–71.
9. M. Sharir and R. Livne, *On minima of functions, intersection patterns of curves and Davenport Schinzel sequences*, Proc. 26th IEEE Symp. on Foundations of Computer Science (1985), 312–320.
10. J. T. Schwartz and M. Sharir, *On the two-dimensional Davenport-Schinzel problem*, J. Symbolic Computation (to appear).
11. R. Wenger, *Upper bounds on geometric permutations for convex sets*, Discrete Comput. Geom. (to appear).

City College, City University of New York, New York, NY 10031
Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, and School of Mathematical Sciences, Tel-Aviv University, 69978 Tel-Aviv, Israel

NYU COMPSCI TR-413
Goodman, Jacob Eli
Common tangents and common
transversals
c.1

NYU COMPSCI TR-413
Goodman, Jacob Eli
Common tangents and common
transversals
c.1

DATE DUE	BORROWER'S NAME

This book may be kept

FOURTEEN DAYS

A fine will be charged for each day the book is kept overtime.

GAYLORD 142			PRINTED IN U S A